Locality-Sensitive Hashing without False Negatives for l_p

Andrzej Pacuk* Piotr Sankowski* Karol Wegrzycki* Piotr Wygocki*

[apacuk,sank,k.wegrzycki,wygos]@mimuw.edu.pl

arXiv:1611.09317v1 [cs.DS] 28 Nov 2016

Abstract

In this paper, we show a construction of locality-sensitive hash functions without false negatives, i.e., which ensure collision for every pair of points within a given radius R in d dimensional space equipped with l_p norm when $p \in [1, \infty]$. Furthermore, we show how to use these hash functions to solve the c-approximate nearest neighbor search problem without false negatives. Namely, if there is a point at distance R, we will certainly report it and points at distance greater than cR will not be reported for $c = \Omega(\sqrt{d}, d^{1-\frac{1}{p}})$. The constructed algorithms work:

- with preprocessing time $\mathcal{O}(n \log(n))$ and sublinear expected query time,
- with preprocessing time $\mathcal{O}(\text{poly}(n))$ and expected query time $\mathcal{O}(\log(n))$.

Our paper reports progress on answering the open problem presented by Pagh [8], who considered the nearest neighbor search without false negatives for the Hamming distance.

1 Introduction

The *Nearest Neighbor* problem is of major importance to a variety of applications in machine learning and pattern recognition. Ordinarily, points are embedded in $\mathbb{R}^d$, and distance metrics usually measure similarity between points. Our task is the following: given a preprocessed set of points $S \subset \mathbb{R}^d$ and a query point $q \in \mathbb{R}^d$, find the point $v \in S$, with the minimal distance to q. Unfortunately, the existence of an efficient algorithm (i.e., whose query and preprocessing time would not depend exponentially on d), would disprove the strong exponential time hypothesis [8, 10]. Due to this fact, we consider the *c-approximate nearest neighbor* problem: given a distance R, a query point q and a constant $c > 1$, we need to find a point within distance cR from point q [4]. This point is called a *cR-near neighbor* of q.

Definition 1. *Point v is an r-near neighbor of q in metric $\mathcal{M}$ iff $\mathcal{M}(q, v) \leq r$.*

One of the most interesting methods for solving the c-approximate nearest neighbor problem in high-dimensional space is *locality-sensitive hashing* (LSH). The algorithm offers a sub-linear query time and a sub-quadratic space complexity. The rudimentary component on which LSH method relies is *locality-sensitive hashing function.* Intuitively, a hash function is *locality-sensitive* if the probability of collision is much higher for "nearby" points than for "far apart" ones. More formally:

Definition 2. *A family $H = \{h : S \to U\}$ is called (r, c, p_1, p_2)-sensitive for distance $\mathcal{D}$ and induced ball $\mathcal{B}(q, r) = \{v : \mathcal{D}(q, v) < r\}$, if for any $v, q \in S$:*

- *if $v \in \mathcal{B}(q, r)$ then $\mathbb{P}[h(q) = h(v)] \geq p_1$,*
- *if $v \notin \mathcal{B}(q, cr)$ then $\mathbb{P}[h(q) = h(v)] \leq p_2$.*

For $p_1 > p_2$ and $c > 1$.

Indyk and Motwani [7] considered randomized c-approximate R-near neighbor (Definition 3).

Definition 3 (The randomized c-approximate R-near neighbor or (R,c)-NN). *Given a set of points in a $P \subset \mathbb{R}^d$ and parameters $R > 0$, $\delta > 0$. Construct a data structure D such that for any query point q, if there exists a R-near neighbor of q in P, D reports some cR-near neighbor of q in P with probability $1 - \delta$.*

*Institute of Informatics, University of Warsaw, Poland

In this paper, we study guarantees for LSH based (R,c)-NN such that for each query point q, every close enough point $\|x-q\|_p < R$ will be certainly returned, i.e., there are no false negatives.[1] In other words, given a set S of size n and a query point q, the result is a set $P \subseteq S$ such that:

$$\{x : \|x-q\|_p < r\} \subseteq P \subseteq \{x : \|x-q\|_p \leq cr\}.$$

Moreover, for each distant point ($\|x-q\|_p > cR$), the probability of being returned is bounded by $\mathsf{p_{fp}}$– probability of false positives. In [8] this type of LSH is called *LSH without false negatives*. The fact that the probability of false negatives is 0 is our main improvement over Indyk and Motwani algorithm [7]. Furthermore, Indyk and Motwani showed that p-stable distributions (where $p \in (0,2]$) are (r, c, p_1, p_2)-sensitive for l_p. We generalized their results on any distribution with mean 0, bounded second and fourth moment and any $p \in [1, \infty]$ (see Lemma 1, for rigorous definitions). Finally, certain distributions from this abundant class guarantee that points within given radius will always be returned (see Figure 1). Unfortunately, our results come with a price, namely $c \geq \max\{\sqrt{d}, d^{1-1/p}\}$.

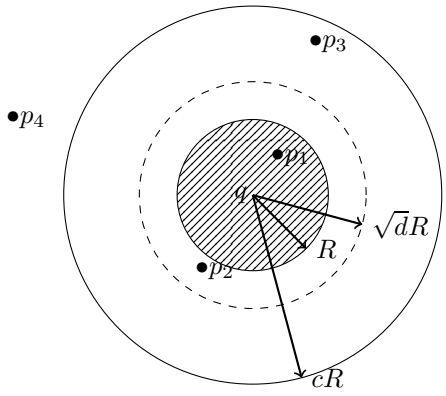

Figure 1: The presented algorithms guarantee that points in the dashed area (p_1) will be reported as neighbors. Points within the dotted circle (p_2) will be reported as neighbor with high probability. Points (p_3) within a distance cR might be reported, but not necessarily. Points (p_4) outside circle cR cannot be reported. The schema picture presents an example for the euclidean distance ($p = 2$).

2 Related Work

2.1 Nearest neighbor in high dimensions

Most common techniques for solving the approximate nearest neighbor search, such as the spatial indexes or k-d trees [3] are designed to work well for the relatively small number of dimensions. The query time for k-d trees is $\mathcal{O}(n^{1-\frac{1}{d}})$ and when the number of dimensions increases the complexity basically converges to $\mathcal{O}(n)$. For interval trees, query time $\mathcal{O}(\log^d n)$ depends exponentially on the number of dimensions. The major breakthrough was the result of Indyk and Motwani [7]. Their algorithm has expected complexity of $\mathcal{O}(dn^{\frac{1}{c}})$ for any approximation constant $c > 1$ and the complexity is tight for any metric l_p (where $p \in (0,2]$). Indyk and Motwani introduced the following LSH functions:

$$h(v) = \left\lfloor \frac{\langle a, v \rangle + b}{r} \right\rfloor,$$

where a is the d-dimensional vector of independent random variables from a p-stable distribution and b is a real number chosen uniformly from the range $[0, r]$.

Our algorithm is based on similar functions and we prove compelling results for more general family of distributions (we show bounds for any distribution with a bounded variance and an expected value equal to 0). Furthermore, our algorithm is correct for any $p \in [1, \infty]$. Indyk and Motwani's LSH algorithm was showed to be optimal for l_1 metric. Subsequently, Andoni et al. [1] showed near optimal results for l_2. Recently, data dependant techniques have been used to further improve LSH by Andoni and Razenshteyn [2]. However, the constant ρ in a query time $\mathcal{O}(n^\rho)$ remains:

$$\rho = \frac{\log p_1}{\log p_2}.$$

When a formal guarantee that $p_1 = 1$ is needed their algorithm does not apply.

[1] $\|\cdot\|_p$ denotes the standard l_p norm for fixed p.

2.2 LSH without false negatives

Recently, Pagh [8] presented a novel approach to nearest neighbor search in Hamming space. He showed the construction of an efficient locality-sensitive hash function family that guarantees collision for any close points. Moreover, Pagh showed that bounds of his algorithm for $cr = \log n/k$ (where $k \in \mathbb{N}$) essentially match bounds of Indyk and Motwani (differ by at most factor $\ln 4$ in the exponent). More precisely, he showed that the problem of false negatives can be avoided in the Hamming space at some cost in efficiency. He proved bounds for general values of c. This paper is an answer to his open problem: whether is it possible to get similar results for other distance measures (e.g., l_1 or l_2).

Pagh introduced the concept of an *r-covering* family of hash function:

Definition 4. *For $A \subseteq \{0,1\}^d$, the Hamming projection family $\mathcal{H}_\mathcal{A}$ is r-covering if for every $x \in \{0,1\}^d$ with $\|x\|_H \leq r$, there exist $h \in \mathcal{H}_\mathcal{A}$ such that $h(x) = \mathbf{0}$.*

Then, he presented a fast method of generating such an r-covering family. Finally, he showed that the expected number of false positives is bounded by $2^{r+1-\|x-y\|_H}$.

3 Basic Construction

We will consider the l_p metric for $p \in [1, \infty]$ and n fixed points in $\mathbb{R}^d$ space. Let v be a d-dimensional vector of independent random variables drawn from distribution $\mathcal{D}$. We define a function h_p as:

$$h_p(x) = \left\lfloor \frac{\langle x, v \rangle}{r\rho_p} \right\rfloor,$$

where $\langle , \rangle$ is a standard inner product and $\rho_p = d^{1-\frac{1}{p}}$. The scaling factor ρ_p is chosen so that: $\|z\|_1 \leq \rho_p \|z\|_p$. The rudimentary distinction between the hash function h_p and LSH is that we consider two hashes equal when they differ at most by one. In Indyk and Motwani [7] version of LSH, there were merely probabilistic guarantees, and close points (say 0.99 and 1.01) could be returned in different buckets with small probability. Since our motivation is to find all close points with absolute certainty, we need to check the adjacent buckets as well.

First, observe that for given points, the probability of choosing a hash function that will classify them as equal is bounded on both sides as given by the following observations. The proofs of these observations are in Appendices A and B.

Observation 1 (Upper bound on the probability of point equivalence)**.**

$$\mathbb{P}\left[|h_p(x) - h_p(y)| \leq 1\right] \leq \mathbb{P}\left[|\langle x - y, v\rangle| < 2\rho_p r\right].$$

Observation 2 (Lower bound on the probability of point equivalence)**.**

$$\mathbb{P}\left[|h_p(x) - h_p(y)| \leq 1\right] \geq \mathbb{P}\left[|\langle x - y, v\rangle| < \rho_p r\right].$$

Interestingly, using the aforementioned observations we can configure a distribution $\mathcal{D}$ so the close points must end up in the same or adjacent bucket.

Observation 3 (Close points have close hashes)**.** *For distribution $\mathcal{D}$ such that every $v_i \sim \mathcal{D}$: $-1 \leq v_i \leq 1$ and for $x, y \in \mathbb{R}^d$, if $\|x - y\|_p < r$ then $\forall_{h_p} |h_p(x) - h_p(y)| \leq 1$.*

Proof. We know that $\|z\|_1 \leq \rho_p \|z\|_p$ and $|v_i| \leq 1$ (because v_i is drawn from bounded distribution $\mathcal{D}$), so:

$$\rho_p \|x - y\|_p \geq \|x - y\|_1 = \sum_i |x_i - y_i| \geq \sum_i |v_i(x_i - y_i)| \quad \geq \left|\textstyle\sum_i v_i(x_i - y_i)\right| = |\langle x - y, v\rangle|.$$

Now, when points are close in l_p:

$$\|x - y\|_p < r \iff \rho_p \|x - y\|_p < \rho_p r \implies |\langle x - y, v\rangle| < \rho_p r.$$

Next, by Observation 2:

$$1 = \mathbb{P}\left[|\langle x - y, v\rangle| < \rho_p r\right] \leq \mathbb{P}\left[|h_p(x) - h_p(y)| \leq 1\right].$$

Hence, the points will inevitably hash into the same or adjacent buckets. □

Now we will introduce the inequality that will help to bound the probability of false positives.

Observation 4 (Inequality of norms in l_p)**.** *Recall that $\rho_p = d^{1-\frac{1}{p}}$. For every $z \in \mathbb{R}^d$ and $p \in [1, \infty]$:*

$$\|z\|_2 \geq \frac{\rho_p}{\max\{d^{\frac{1}{2}}, d^{1-\frac{1}{p}}\}} \|z\|_p.$$

This technical observation is proven in Appendix C.

The major question arises: what is the probability of false positives? In contrast to the Indyk and Motwani [7], we cannot use p-stable distributions because these distributions are not bounded. We will present the proof for a different class of functions.

Lemma 1 (The probability of false positives for general distribution)**.** *Let $\mathcal{D}$ be a random variable such that $\mathbb{E}(\mathcal{D}) = 0$, $\mathbb{E}(\mathcal{D}^2) = \alpha^2$, $\mathbb{E}(\mathcal{D}^4) \leq 3\alpha^4$ (for any $\alpha \in \mathbb{R}^+$). Define constant $\tau_1 = \frac{2}{\alpha} \max\{d^{\frac{1}{2}}, d^{1-\frac{1}{p}}\}$.*

When $\|x - y\|_p > cr$, $x, y \in \mathbb{R}^d$ and $c > \tau_1$ then:

$$\mathsf{p}_{\mathsf{fp}_1} = \mathbb{P}\left[|h_p(x) - h_p(y)| \leq 1\right] < 1 - \frac{\left(1 - \frac{\tau_1^2}{c^2}\right)^2}{3},$$

for every metric l_p, where $p \in [1, \infty]$ ($\mathsf{p}_{\mathsf{fp}_1}$ is the probability of false positive).

Proof. By Observation 4:

$$\|z\|_2 \geq \frac{2\|z\|_p}{\alpha\tau_1}\rho_p$$

Subsequently, let $z = x - y$ and define a random variable $X = \langle z, v\rangle$. Therefore:

$$\mathbb{E}(X^2) = \alpha^2\|z\|_2^2 \geq (\frac{2\|z\|_p}{\tau_1}\rho_p)^2 > (2r\rho_p\frac{c}{\tau_1})^2.$$

Because $\frac{c}{\tau_1} > 1$ we have $\theta = \frac{(2r\rho_p)^2}{\mathbb{E}X^2} < 1$. Variable θ and a random variable $X^2 > 0$ satisfy Paley-Zygmunt inequality (analogously to [9]):

$$\begin{aligned}\mathbb{P}\left[|h_p(x) - h_p(y)| > 1\right] &\geq \mathbb{P}\left[|\langle z, v\rangle| \geq 2r\rho_p\right] \geq \mathbb{P}\left[X^2 > (2r\rho_p)^2\right] \\ &\geq \left(1 - \frac{(2r\rho_p)^2}{\mathbb{E}(X^2)}\right)^2 \frac{\mathbb{E}(X^2)^2}{\mathbb{E}(X^4)}.\end{aligned}$$

Eventually, we assumed that $\mathbb{E}(X^4) \leq 3(\alpha\|z\|_2)^4$:

$$\mathbb{P}\left[|h_p(x) - h_p(y)| > 1\right] \geq \frac{\left(1 - \frac{(2r\rho_p)^2}{\mathbb{E}(X^2)}\right)^2}{3} > \frac{\left(1 - \frac{\tau_1^2}{c^2}\right)^2}{3}.$$

□

Simple example of a distribution that satisfies both Observation 3 and Lemma 1 is a uniform distribution on $(-1, 1)$ with a standard deviation α equal to $\sqrt{\frac{1}{3}}$. Another example of such distribution is a discrete distribution with uniform values $\{-1, 1\}$. As it turns out, Lemma 2 shows that the discrete distribution leads to even better bounds.

Lemma 2 (Probability of false positives for the discrete distribution)**.** *Let $\mathcal{D}$ be a random variable such that $\mathbb{P}\left[\mathcal{D} = \pm 1\right] = \frac{1}{2}$. Define constant $\tau_2 = \sqrt{8}\max\{d^{\frac{1}{2}}, d^{1-\frac{1}{p}}\}$. Then for every $p \in [1, \infty]$, $x, y \in \mathbb{R}^d$ and $c > \tau_2$ such that $\|x - y\|_p > cr$, it holds:*

$$\mathsf{p}_{\mathsf{fp}_2} = \mathbb{P}\left[|h_p(x) - h_p(y)| \leq 1\right] < 1 - \frac{(1 - \frac{\tau_2}{c})^2}{2}.$$

Proof. Because of Observation 4 we have the inequality:

$$\|z\|_2 \geq \sqrt{8}\frac{\|z\|_p}{\tau_2}\rho_p.$$

Let $z = x - y$ and $X = \langle z, v\rangle$, be a random variable. Then:

$$\mathbb{P}\left[|h_p(x) - h_p(y)| > 1\right] \geq \mathbb{P}\left[|X| > 2r\rho_p\right].$$

Khintchine inequality [5] states $\mathbb{E}|X| \geq \frac{\|z\|_2}{\sqrt{2}}$, so:

$$\mathbb{E}(|X|) \geq \frac{\|z\|_2}{\sqrt{2}} \geq \frac{2\rho_p \|z\|_p}{\tau_2} > 2r\rho_p \frac{c}{\tau_2}.$$

Note that, a random variable $|X|$ and $\theta = \frac{2r\rho_p}{\mathbb{E}(|X|)} < 1$, satisfy the Paley-Zygmunt inequality (because $\frac{c}{\tau_2} > 1$), though:

$$\begin{aligned}\mathbb{P}\left[h_p(x) - h_p(y)| > 1\right] &\geq \left(1 - \frac{2r\rho_p}{\mathbb{E}(|X|)}\right)^2 \frac{\mathbb{E}(|X|)^2}{\mathbb{E}(|X|^2)} \\ &> \left(1 - \frac{2r\rho_p}{2r\rho_p \frac{c}{\tau_2}}\right)^2 \frac{1}{2} = \frac{\left(1 - \frac{\tau_2}{c}\right)^2}{2}.\end{aligned}$$

□

Altogether, in this section we have introduced a family of hash functions h_p which:

- guarantees that, with an absolute certainty, points within the distance R will be mapped to the same or adjacent buckets (see Observation 3),
- maps "far away" points to the non-adjacent hashes with high probability (Lemma 1 and Lemma 2).

These properties will enable us to construct an efficient algorithm for solving the c-approximate nearest neighbor search problem without false negatives.

3.1 Tightness of bounds

We showed that for two distant points $x, y : \|x - y\|_p > cr$, the probability of a collision is small when $c = \max\{\rho_p, \sqrt{d}\}$. The natural question arises: Can we bound the probability of a collision for points $\|x-y\|_p > c'r$ for some $c' < c$?

We will show that such c' does not exist, i.e., there always exists $\tilde{x}$ such that $\|\tilde{x}\|_p$ will be arbitrarily close to cr, so $\tilde{x}$ and $\vec{0}$ will end up in the same or adjacent bucket with high probability. More formally, for any $p \in [1, \infty]$, for $h_p(x) = \left\lfloor \frac{\langle x,v \rangle}{r\rho_p} \right\rfloor$, where coordinates of d-dimensional vector v are random variables v_i, such that $-1 \leq v_i \leq 1$ with $\mathbb{E}(v_i) = 0$. We will show that there always exists $\tilde{x}$ such that $\|\tilde{x}\|_p \approx r \max\{\rho_p, \sqrt{d}\}$ and $|h_p(\tilde{x}) - h_p(\vec{0})| \leq 1$ with high probability.

For $p \geq 2$ denote $x_0 = (r\rho_p - \epsilon, 0, 0, \ldots, 0)$. We have $\|x_0 - \vec{0}\|_p = r\rho_p - \epsilon$ and:

$$|h_p(x_0) - h_p(\vec{0})| = \left| \left\lfloor \frac{r\rho_p - \epsilon}{r\rho_p} \cdot v_1 \right\rfloor - 0 \right| \leq 1.$$

For $p \in [1, 2)$, denote $x_1 = rd^{-\frac{1}{p}+\frac{1}{2}-\epsilon}\vec{1}$. We have $\|x_1\|_p = rd^{\frac{1}{2}-\epsilon}$ and by applying Observation 2 for complementary probabilities:

$$\begin{aligned}\mathbb{P}\left[|h_p(x_1) - h_p(\vec{0})| > 1\right] &\leq \mathbb{P}\left[|\langle x_1, v\rangle| \geq \rho_p r\right] = \mathbb{P}\left[|\langle \vec{1}, v\rangle| \geq d^{\frac{1}{2}+\epsilon}\right] \\ &= \mathbb{P}\left[\left|\frac{\sum_{i=1}^d v_i}{d}\right| \geq d^{-\frac{1}{2}+\epsilon}\right] \leq 2 \cdot \exp\left(\frac{-d^{2\epsilon}}{2}\right).\end{aligned}$$

The last inequality follows from Hoeffding [6] (see Appendix D for technical details).

So the aforementioned probability for $p \in [1, 2)$ is bounded by an expression exponential in $d^{2\epsilon}$. Even if we would concatenate k random hash functions (see proof of Theorem 1 for more details), the chance of collision would be at least $(1 - 2e^{\frac{-d^{2\epsilon}}{2}})^k$. To bound this probability, the number k needs to be at least $\Theta(e^{\frac{d^{2\epsilon}}{2}})$. The probability bounds do not work for ϵ arbitrary close to 0: we proved that introduced hash functions for $c = d^{1/2-\epsilon}$ do not work (may give false positives).[2]

Hence, to obtain a significantly better approximation factor c, one must introduce a completely new family of hash functions.

[2]However, one may try to obtain tighter bound (e.g., $c = d^{1/2}/\log(d)$) or show that for every $\epsilon > 0$, the approximation factor $c = d^{1/2} - \epsilon$ does not work.

4 The algorithm

In this section, we apply the LSH family introduced in Section 3 to construct an *c-approximate* algorithm without false negatives. To begin with, we will define a general algorithm that will satisfy our conditions. Subsequently, we will show that complexity of the query is sublinear, and it depends linearly on the number of dimensions.

Theorem 1. *For any $c > \tau$ and the number of iterations $k \geq 0$, there exists a c-approximate nearest neighbor algorithm without false negatives for l_p, where $p \in [1, \infty]$:*

- *Preprocessing time:* $\mathcal{O}(n(kd + 3^k))$,
- *Memory usage:* $\mathcal{O}(n3^k)$,
- *Expected query time:* $\mathcal{O}(d(|P| + k + n\mathsf{p_{fp}}^k))$.

Where $|P|$ is the size of the result and $\mathsf{p_{fp}}$ is the upper bound of probability of false positives (note that $\mathsf{p_{fp}}$ depends on a choice of τ from Lemma 1 or Lemma 2).

Proof. Let $g(x) \coloneqq (h_p^1(x), h_p^2(x), \ldots, h_p^k(x))$ be a hash function defined as a concatenation of k random LSH functions presented in Section 3. We introduce the clustering $m : g(\mathbb{R}^d) \to 2^n$, where each cluster is assigned to the corresponding hash value. For each hash value α, the corresponding cluster $m(\alpha)$ is $\{x : g(x) = \alpha\}$.

Since we consider two hashes to be equal when they differ at most by one (see Observation 3), for hash α, we need to store the reference for every point that satisfies $\|\alpha - x\| \leq 1$. The number of such clusters is 3^k, because the result of each hash function can vary by one of $\{-1, 0, 1\}$ and the number of hash functions is k. Thus, the memory usage is $\mathcal{O}(n3^k)$ (see Figure 2).

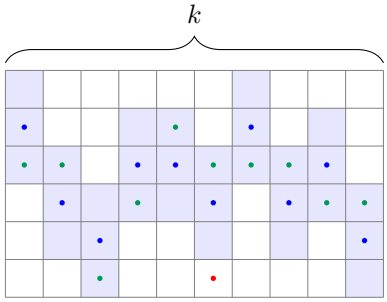

Figure 2: Blue dots represent value of $g(q)$ for query. Green dots are always distant by 1, hence green and blue points are considered close. At least one red dot is distant from blue dot by more than 1, hence red dots will not be considered close to blue. Thus, algorithm needs to check 3^k various possibilities.

To preprocess the data, we need to compute the value of the function g for every point in the set and then put its reference into 3^k cells. Hence, the preprocessing time complexity equals $\mathcal{O}(n(kd + 3^k))$.

Eventually, to answer a query, we need to compute $g(q)$ in time $\mathcal{O}(kd)$ and then for every point in $\|g(x) - g(q)\|_\infty \leq 1$ remove distant points $\|x - q\|_p > cR$. Hence, we need to look up every false-positive to check whether they are within distance cr from the query point. We do that in expected time $\mathcal{O}(d(|P| + k + n\mathsf{p_{fp}}^k))$, because $n\mathsf{p_{fp}}^k$ is the expected number of false positives. □

The number of iterations k can be chosen arbitrarily, so we will choose the optimal value to minimize the query time. This gives the main result of this paper:

Theorem 2. *For any $c > \tau$ and for large enough n, there exists a c-approximate nearest neighbor algorithm without false negatives for l_p, where $p \in [1, \infty]$:*

- *Preprocessing time:* $\mathcal{O}(n(\gamma d \log n + (\frac{n}{d})^\gamma)) = \text{poly}(n)$,
- *Memory usage:* $\mathcal{O}(n(\frac{n}{d})^\gamma)$,
- *Expected query time:* $\mathcal{O}(d(|P| + \gamma \log(n) + \gamma d))$.

Where $|P|$ is the size of the result, $\gamma = \frac{\ln 3}{-\ln \mathsf{p_{fp}}}$ and $\mathsf{p_{fp}}$ and τ are chosen as in Theorem 1.

Proof. Denote $a = -\ln \mathsf{p_{fp}}$, $b = \ln 3$ and set $k = \left\lceil \frac{\ln \frac{na}{d}}{a} \right\rceil$.

Let us assume that n is large enough so that $k \geq 1$. Then using the fact that $x^{1/x}$ is bounded for $x > 0$ we have:

$$3^k \leq 3 \cdot (3^{\ln \frac{na}{d}})^{1/a} = 3 \cdot (\frac{na}{d})^{b/a} = \mathcal{O}((\frac{n}{d})^{b/a}) = \mathcal{O}((\frac{n}{d})^{\gamma}),$$

$$n{\mathsf{p_{fp}}}^k = ne^{-ak} \leq ne^{-a\frac{\ln(\frac{na}{d})}{a}} = \frac{d}{a} = \mathcal{O}(d\gamma),$$

$$k = \mathcal{O}(\gamma \log(n)).$$

Substituting these values in the Theorem 1 gives needed complexity guaranties. □

There are two variants of Theorems 1, 2 and 3. In the first variant, we show complexity bounds for very general class of hashing functions introduced in Lemma 1. In the second one, we show slightly better guaranties for hashing functions which are generated using discrete probability distribution on $\{0, 1\}$ introduced in Lemma 2. For simplicity the following discussion is restricted only to the second variant which gives better complexity guaranties. The definitions of constants $\mathsf{p_{fp}}_2$ and τ_2 used in this discussion are taken from Lemma 2 . For a general case, i.e., $\mathsf{p_{fp}}_1$ and τ_1 taken from Lemma 1, we get only slightly worse results.

The complexity bounds introduced in the Theorem 2 can be simplified using the fact that $\ln(x) < x - 1$. Namely, we have:

$$\gamma = \frac{\ln 3}{-\ln \mathsf{p_{fp}}_2} = \frac{\ln 3}{-\ln(1 - \frac{(1-\frac{\tau_2}{c})^2}{2})} < \frac{2 \ln 3}{(1 - \frac{\tau_2}{c})^2}.$$

However, the preprocessing time is polynomial in n for any constant c, it strongly depends on the bound for probability $\mathsf{p_{fp}}_2$ and c. Particularly when c is getting close to τ_2, the exponent of the preprocessing time might be arbitrarily large.

To the best of our knowledge, this is the first algorithm that will ensure that no false negatives will be returned by the nearest neighbor approximated search and does not depend exponentially on the number of dimensions. Note that for given c, the parameter γ is fixed. By Lemma 2, we have: $\mathsf{p_{fp}}_2 = 1 - \frac{(1-\frac{\tau_2^2}{c^2})^2}{2}$, so:

$$\lim_{c \to \infty} \gamma = \lim_{c \to \infty} \frac{\ln 3}{-\ln \mathsf{p_{fp}}_2} = \log_2 3 \approx 1.58.$$

If we omit terms polynomial in d, the preprocessing time of the algorithm from Theorem 2 converges to $\mathcal{O}(n^{2.58})$ ($\mathcal{O}(n^{3.71})$ for general case - see Appendix E.).

4.1 Light preprocessing

Although the preprocessing time $\mathcal{O}(n^{2.58})$ may be reasonable when there are multiple, distinct queries and the data set does not change (e.g., static databases, pre-trained classification, geographical map). Still, unless the number of points is small, this algorithm does not apply. Here, we will present an algorithm with a light preprocessing time $\mathcal{O}(dn \log n)$ and $\mathcal{O}(n \log n)$ memory usage where the expected query time is $o(n)$.

The algorithm with light preprocessing is very similar to the algorithm described in Theorem 1, but instead of storing references to the point in all 3^k buckets during preprocessing, this time searching for every point x that matches $\|x - q\|_\infty \leq 1$ is done during the query time.

The expected query time with respect to k is $\mathcal{O}(d(|P| + k + n{\mathsf{p_{fp}}}^k) + 3^k)$. During the preprocessing phase we only need to compute k hash values for each of n points and store them in memory. Hence, preprocessing requires $\mathcal{O}(kdn)$ time and uses $\mathcal{O}(nk)$ memory.

Theorem 3. *For any $c > \tau$ and for large enough n, there exists a c-approximate nearest neighbor algorithm without false negatives for l_p, where $p \in [1, \infty]$:*

- *Preprocessing time: $\mathcal{O}(nd \log n)$,*
- *Memory usage: $\mathcal{O}(n \log n)$,*
- *Expected query time: $\mathcal{O}(d(|P| + n^{\frac{b}{a+b}}(\frac{b}{a})^{\frac{a}{b+a}}))$.*

Where $|P|$ is the size of the result, $a = -\ln \mathsf{p_{fp}}$, $b = \ln 3$, $\mathsf{p_{fp}}$ and τ are chosen as in Theorem 1.

Proof. We set the number of iterations $k = \left\lceil \frac{\ln \frac{na}{b}}{a+b} \right\rceil$. Assume n needs to be large enough so that $k \geq 1$. Since a is upper bounded for both choices of $\mathsf{p_{fp}}$:

$$3^k \leq 3 \cdot 3^{\frac{\ln(\frac{na}{b})}{a+b}} = 3(\frac{na}{b})^{\frac{b}{a+b}} = \mathcal{O}(n^{\frac{b}{a+b}}).$$

Analogously:

$$n{\mathsf{p_{fp}}}^k = n(e^{-a})^k \leq ne^{-a\frac{\ln(\frac{na}{b})}{a+b}} = n \cdot \left(\frac{b}{a}\right)^{\frac{a}{a+b}} \cdot \left(\frac{1}{n}\right)^{\frac{a}{a+b}} = n^{\frac{b}{a+b}} \left(\frac{b}{a}\right)^{\frac{a}{a+b}}.$$

Hence, for this choice of k we obtain the expected query time is equal to:

$$\begin{aligned}\mathcal{O}(d(|P| + k + n{\mathsf{p_{fp}}}^k)) + 3^k &= \mathcal{O}(d(|P| + \log n + n^{\frac{b}{a+b}} \left(\frac{b}{a}\right)^{\frac{a}{a+b}}) + n^{\frac{b}{a+b}}) \\ &= \mathcal{O}(d(|P| + n^{\frac{b}{a+b}} \left(\frac{b}{a}\right)^{\frac{a}{a+b}}).\end{aligned}$$

Substituting k, we obtain formulas for preprocessing time and memory usage.

□

Eventually, exactly as previously for a general distribution from Lemma 1, when $c \to \infty$ we have: $a \to \ln \frac{3}{2}$ (see Theorem 3 for the definition of constant a). Hence, for a general distribution we have a bound for complexity equal to $\mathcal{O}(n^{\log_{4.5} 3}) \approx \mathcal{O}(n^{0.73})$. For the discrete distribution from Lemma 2, the constant a converges to $\ln 2$. Hence, the expected query time converges to $\mathcal{O}(n^{0.61})$.

5 Conclusion and Future Work

We have presented the c-approximate nearest neighbor algorithm without false negatives in l_p for all $p \in [1, \infty]$ and $c > \max\{\sqrt{d}, d^{1-1/p}\}$. Due to this inequality our algorithm can be used cognately to the original LSH [7] but with additional guarantees about very close points (one can set $R' = \sqrt{d}R$ and be certain that all points within distance R will be returned). In contrast to the original LSH, our algorithm does not require any additional parameter tunning.

The future work concerns relaxing restriction on the approximation factor c and reducing time complexity of the algorithm or proving that these restrictions are essential. We wish to match the time complexities given by [7] or show that achieved bounds are optimal. We show the tightness of our construction, hence to break the bound of $\sqrt{d}$, one would need to introduce a new technique.

6 Acknowledgments

This work was supported by ERC PoC project PAAl-POC 680912 and FET project MULTIPLEX 317532. We would also like to thank Rafał Latała for meaningful discussions.

References

[1] Alexandr Andoni and Piotr Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Commun. ACM*, 51(1):117–122, 2008.

[2] Alexandr Andoni and Ilya Razenshteyn. Optimal data-dependent hashing for approximate near neighbors. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 793–801. ACM, 2015.

[3] Jon Louis Bentley. K-d trees for semidynamic point sets. In *Proceedings of the Sixth Annual Symposium on Computational Geometry*, SCG '90, pages 187–197, New York, NY, USA, 1990. ACM.

[4] Mayur Datar and Piotr Indyk. Locality-sensitive hashing scheme based on p-stable distributions. In *In SCG 04: Proceedings of the Twentieth Annual Symposium on Computational Geometry*, pages 253–262. ACM Press, 2004.

[5] Uffe Haagerup. The best constants in the Khintchine inequality. *Studia Mathematica*, 70(3):231–283, 1981.

[6] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, March 1963.

[7] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, pages 604–613, New York, NY, USA, 1998. ACM.

[8] Rasmus Pagh. Locality-sensitive hashing without false negatives. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1–9. SIAM, 2016.

[9] Mark Veraar. On Khintchine inequalities with a weight. *Proceedings of the American Mathematical Society*, 138:4119–4121, 2010.

[10] Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2):357–365, December 2005.

A Proof of Observation 1

Proof. We will use, the fact that for any $x, y \in \mathbb{R}$ we have $|\lfloor x \rfloor - \lfloor y \rfloor| \leq 1 \Rightarrow |x - y| < 2$. Then the following implications hold:

$$|h_p(x) - h_p(y)| \leq 1 \iff \left| \left\lfloor \frac{\langle x,v \rangle}{\rho_p r} \right\rfloor - \left\lfloor \frac{\langle y,v \rangle}{\rho_p r} \right\rfloor \right| \leq 1 \implies \left| \frac{\langle x,y \rangle}{\rho_p r} - \frac{\langle y,v \rangle}{\rho_p r} \right| < 2 \iff$$
$$\iff \quad |\langle x - y, v \rangle| < 2\rho_p r.$$

So, based on the increasing property of the probability:

$$\text{if } A \subset B \text{ then } \mathbb{P}[A] \leq \mathbb{P}[B],$$

the inequality of the probabilities holds.

□

B Proof of Observation 2

Proof. We will use the fact that for $x, y \in \mathbb{R} : |x - y| < 1 \Rightarrow |\lfloor x \rfloor - \lfloor y \rfloor| \leq 1$).

$$\left| \langle x - y, v \rangle \right| < \rho_p r \iff \left| \frac{\langle x,v \rangle}{\rho_p r} - \frac{\langle x,v \rangle}{\rho_p r} \right| < 1 \implies \left| \left\lfloor \frac{\langle x,v \rangle}{\rho_p r} \right\rfloor - \left\lfloor \frac{\langle x,v \rangle}{\rho_p r} \right\rfloor \right| \leq 1 \iff$$
$$\iff \quad |h_p(x) - h_p(y)| \leq 1$$

□

C Proof of Observation 4

Proof. For every $0 < b \leq a$ vectors in $\mathbb{R}^d$ satisfy the inequality:

$$\|z\|_a \leq \|z\|_b \leq d^{(\frac{1}{b} - \frac{1}{a})} \|z\|_a. \tag{1}$$

For $p > 2$ we have $\max\{d^{\frac{1}{2}}, d^{1-\frac{1}{p}}\} = d^{1-\frac{1}{p}}$. Then, using ineqaulity (1) for $a = p$ and $b = 2$ we have:

$$\|z\|_2 \geq \|z\|_p = \frac{\rho_p}{d^{1-\frac{1}{p}}} \|z\|_p = \frac{\rho_p}{\max\{d^{\frac{1}{2}}, d^{1-\frac{1}{p}}\}} \|z\|_p$$

For $1 \leq p \leq 2$ we have $\max\{d^{\frac{1}{2}}, d^{1-\frac{1}{p}}\} = d^{\frac{1}{2}}$. Analogously by using inequality (1) for $a = 2$ and $b = p$:

$$\|z\|_p \leq d^{\frac{1}{p}-\frac{1}{2}} \|z\|_2 = \|z\|_2 \frac{d^{\frac{1}{2}}}{\rho_p}$$

Hence, by dividing both sides we have:

$$\|z\|_p \frac{\rho_p}{\max\{d^{\frac{1}{2}}, d^{1-\frac{1}{p}}\}} \leq \|z\|_2$$

□

D Hoeffding bound

Here we are going to show all technical details used in the proof in the Section 3.1. Let us start with the Hoeffding inequality. Let $X_1, \ldots, X_d$ be bounded independent random variables: $a_i \leq X_i \leq b_i$ and $\overline{X}$ be the mean of these variables $\overline{X} = \sum_{i=1}^{d} X_i / d$. Theorem 2 of Hoeffding [6] states:

$$\mathbb{P}\left[|\overline{X} - \mathbb{E}\left[\overline{X}\right]| \geq t \right] \leq 2 \cdot \exp\left(- \frac{2d^2 t^2}{\sum_{i=1}^{d} (b_i - a_i)^2} \right).$$

In our case, $D_1, \ldots, D_d$ are bounded by $a_i = -1 \leq D_i \leq 1 = b_i$ with $\mathbb{E}D_i = 0$. Hoeffding inequality implies:

$$\mathbb{P}\left[\left|\frac{\sum_{i=1}^{d} D_i}{d}\right| \geq t\right] \leq 2 \cdot \exp\left(-\frac{2d^2t^2}{\sum_{i=1}^{d}(b_i - a_i)^2}\right) = 2 \cdot \exp\left(-\frac{dt^2}{2}\right).$$

Taking $t = d^{-1/2+\epsilon}$ we get the claim:

$$\mathbb{P}\left[\left|\frac{\sum_{i=1}^{d} D_i}{d}\right| \geq d^{-1/2+\epsilon}\right] \leq 2 \cdot \exp\left(-\frac{d^{2\epsilon}}{2}\right).$$

E Preprocessing complexity bounds for the distributions introduced in Lemma 1

By Lemma 1, we have: $\mathsf{p}_{\mathsf{fp}_1} = 1 - \frac{(1-\frac{\tau_1^2}{c^2})^2}{3}$, so:

$$\lim_{c\to\infty} \gamma = \lim_{c\to\infty} \frac{\ln 3}{-\ln \mathsf{p}_{\mathsf{fp}_1}} = \frac{\ln 3}{\ln 1.5} \approx 2.71.$$

If we omit terms polynomial in d, the preprocessing time of the algorithm from Theorem 2 converges to $\mathcal{O}(n^{3.71})$.